arXiv:1605.07844v2 [cs.IR] 8 Oct 2016

Dimension Projection among Languages based on Pseudo-relevant Documents for Query Translation

Javid Dadashkarimi, Mahsa S. Shahshahani, Amirhossein Tebbifakhr, Heshaam Faili, and Azadeh Shakery
School of ECE, College of Engineering, University of Tehran, Iran
{dadashkarimi,ms.shahshahani,a.tebbifakhr,hfaili,shakery}@ut.ac.ir

Abstract. Using top-ranked documents in response to a query has been shown to be an effective approach to improve the quality of query translation in dictionary-based cross-language information retrieval. In this paper, we propose a new method for dictionary-based query translation based on dimension projection of embedded vectors from the pseudo-relevant documents in the source language to their equivalents in the target language. To this end, first we learn low-dimensional vectors of the words in the pseudo-relevant collections separately and then aim to find a query-dependent transformation matrix between the vectors of translation pairs appeared in the collections. At the next step, representation of each query term is projected to the target language and then, after using a softmax function, a query-dependent translation model is built. Finally, the model is used for query translation. Our experiments on four CLEF collections in French, Spanish, German, and Italian demonstrate that the proposed method outperforms a word embedding baseline based on bilingual shuffling and a further number of competitive baselines. The proposed method reaches up to 87% performance of machine translation (MT) in short queries and considerable improvements in verbose queries.

1 Introduction

Pseudo-relevance feedback (PRF) has long been shown to be an effective approach for updating query language models in information retrieval (IR) [6,12,4,3]. In cross-language environments where there are a couple of document sets in different languages, building such models requires bridging the gap of the languages. To this end, cross-lingual topical relevance models (CLTRLM) aims to find a way to transform knowledge of the sets to the query model using bilingual topic modeling and a bilingual dictionary [3,1]. The performance of CLTRLM is heavily depends on the number of alignments in a comparable corpora and their qualities as well. Recently, bilingual word embedding is tailored effectively to this end where low-dimensional vectors are built after shuffling all the alignments [10]. However the effectiveness of this method has not been investigated at cross-lingaul PRF yet.

In this paper we propose a new method for building translation models on pseudo-relevant collections using a neural network-based language model. The proposed cross-lingual word embedding translation model (CLWETM) takes advantage of a query-dependent transformation matrix between low-dimensional vectors of the languages. Indeed, we aim to find a transformation matrix to bring the vector of each query term, built on the source collection, to dimensionality of the target language and then compute the translation probabilities based on a softmax function. To this aim, first we learn word representations of the pseudo-relevant collections separately and then focus on finding a transformation matrix minimizing a distance function between all translation pairs appeared at the collections. This method captures semantics of both the collections with a rotation and a scaling

embedded in the matrix. Finally, a softmax function is used to build a query-dependent translation model based on similarity of transformed vector of each query term with the vectors of its translation candidates in the target language.

Unlike CLTRLM and the mixed word embedding translation model (MIXWETM) based on shuffling alignments in a comparable corpora ([10]), CLWETM considers quite sentence-level contexts of the words and therefore captures deeper levels of n-grams in both the languages. Furthermore, the obtained model can be incorporated within a language modeling framework, the state-of-the-art retrieval framework in the literature, and therefore the proposed method does not suffer from disadvantages of low-dimensional query-document similarity in ad-hoc retrieval [10].

Experimental results on four CLEF collections in French, German, Spanish, and Italian demonstrate that the proposed method outperforms all competitive baselines of dictionary-based cross-language information retrieval (CLIR) in language modeling when it is combined with a global translation model. The proposed method reaches up to 83% performance of the monolingual run and 87% performance of machine translation in short queries. CLWETM has more better results in verbose queries and even improvements compared to MT in the Italian collection.

2 Related Works

Pseudo-relevance feedback has long been employed as a powerful method for estimating query language models in a large number of studies [6,5,12]. Cross-lingual relevance model (CLRLM) and CLTRLM are state-of-the-art methods in cross-lingual environments [1,11,3]. Unlike CLRLM that depends on parallel corpora and bilingual lexicons, CLTRLM aims at finding a number of bilingual topical variables from a comparable corpus in order to transfer relevance score of a term from one language to another. Ganguly et al. proposed to use these variables for query translation and demonstrated that CLTRLM is an effective method particularly for resource-lean languages [1].

In CLTRLM top-ranked documents $F^s = \{d_1^s, d_2^s, .., d_{|F^s|}^s\}$ retrieved in response to the source query ($\mathbf{q}^s$) and top-ranked documents $F^t = \{d_1^t, d_2^t, .., d_{|F^t|}^t\}$ retrieved in response to a translation of the query ($\mathbf{q}^t$) are assumed to be relevant documents and then we expect that each word w^t in target language is generated either from a target event or a source event as follows: $p(w^t|\mathbf{q}^s) = p(w^t|\mathbf{z}^t)p(\mathbf{z}^t|\mathbf{q}^s) + p(w^t|\mathbf{w}^s)p(\mathbf{w}^s|\mathbf{q}^s)$ in which $\mathbf{z}^t$ is a topical variable on F^t and $\mathbf{w}^s$ is a translation of w^t in the dictionary (see whole the details in [1]).

It is noteworthy that topic modeling considers co-occurrences of the terms within documents without considering their sentence-level contexts. Language modeling based on neural networks is a popular technique for capturing co-occurrences of the terms within a constant window c and thus, this model embeds semantic of the language as well as deeper levels of n-grams [8].

But when it comes to cross-lingual environments a constraint is required for integrating dimensions in the languages. There are a couple of methods to this end; on-line methods and off-line methods. On-line methods aim at finding a unified space for the languages during the learning process as follows [2]: $\mathcal{L}(\theta) = \mathcal{L}(\theta^s) + \mathcal{L}(\theta^t) + \lambda f(\theta^s, \theta^t)$ where the last term is a regularization term. Recently, Vulic et al., introduced the shuffling-based word embedding method over comparable corpora in which the alignments play as a number of constraints on the vectors [10]. Ultimately, this approach considers neither the absolute/relative positions of the terms nor their sentence level contexts for estimation.

Off-line methods learn the vectors separately and then find a transformation matrix minimizing a constraint function f as follows [7]: $f(\mathbf{W}) = \sum_{\mathbf{x},\mathbf{z}} ||\mathbf{W}^T\mathbf{x} - \mathbf{z}||^2$ in which

$\mathbf{x} \in \mathbb{R}^{n\times 1}$ and $\mathbf{z} \in \mathbb{R}^{n\times 1}$ are the vectors of translation pairs from a dictionary or a parallel corpus.

3 Linear Projection between Languages based on Pseudo-relevant Documents

In this section we introduce the proposed method in more details. We employ an off-line approach for learning bilingual representations of the words by exploiting pseudo-relevant documents in both source and target languages. To this end, first we learn word representations of the pseudo-relevant collections separately and then focus on finding a transformation matrix minimizing a distance function between all translation pairs appeared on the collections. As shown in Equation 1 our goal is to minimize f with respect to a transformation matrix $\mathbf{W} \in \mathbb{R}^{n\times n}$; f is defined as follow:

$$f(\mathbf{W}) = \sum_{(w^s, w^t)} \frac{1}{2} ||\mathbf{W}^T \mathbf{u}_{w^s} - \mathbf{v}_{w^t}||^2 \tag{1}$$

where, w^t is a translation pair of w^s from F^s that is appeared in F^t; $\mathbf{u}_{w^s} \in \mathbb{R}^{n\times 1}$ and $\mathbf{u}_{w^t} \in \mathbb{R}^{n\times 1}$ are the corresponding vectors respectively. To solve this problem we choose the stochastic gradient descent algorithm (i.e., $\frac{\partial f}{\partial \mathbf{W}} = 0$):

$$\mathbf{W}^{t+1} = \mathbf{W}^t - \eta(\mathbf{W}^T \mathbf{u}_{w^s} - \mathbf{v}_{w^t})\mathbf{u}_{w^s}^T \tag{2}$$

where η is a constant learning rate. $\mathbf{W}$ is to be initialized randomly and then be updated incrementally.

3.1 Bilingual Representations and Translation Models

In this section we introduce a method based on bilingual word representations for building a translation model for query and then incorporate it within language modeling, the state-of-the-art retrieval framework.

$$\hat{\mathbf{u}}_{w^s} = \mathbf{W}^T \mathbf{u}_{w^s} \tag{3}$$

where $\mathbf{W}^T \mathbf{u}_{w^s}$ transforms the source vector built on the source collection to the target low-dimensional space. The new translation model is built as follows:

$$p(w_t|w_s) = \frac{e^{\frac{\hat{\mathbf{u}}_{w^s} \cdot \mathbf{v}_{w^t}}{||\hat{\mathbf{u}}_{w^s}|| \; ||\mathbf{v}_{w^t}||}}}{\sum_{\bar{w}^t \in \mathcal{T}\{w^s\}} e^{\frac{\hat{\mathbf{u}}_{w^s} \cdot \mathbf{v}_{\bar{w}^t}}{||\hat{\mathbf{u}}_{w^s}|| \; ||\mathbf{v}_{\bar{w}^t}||}}} \tag{4}$$

where $\mathcal{T}\{w^s\}$ is the list of translations of w^s. Instead of topical information propagation taking place on CLTRLM and joint cross-lingual topical relevance model (JCLTRLM), CLWETM tailors semantic projection and scaling both embedded in $\mathbf{W}$. In other words, CLWETM embeds quite sentence-level contexts of the words and thus captures deeper levels of n-grams as well.

Combining Translation Models Since the obtained model is a probabilistic translation model we can interpolate it with other models using a constant controlling parameter as follow:

$$p(w^t|w^s) = \alpha p_1(w^t|w^s) + (1-\alpha) p_2(w^t|w^s) \tag{5}$$

Table 1: Collection Characteristics

ID	Lang.	Collection	Queries	#docs	#qrels
IT	Italy	La Stampa 94, AGZ 94	CLEF 2003-2003, Q:91-140	108,577	4,327
SP	Spanish	EFE 1994	CLEF 2002, Q:91-140	215,738	1,039
DE	German	Frankfurter Rundschau 94, SDA 94, Der Spiegel 94-95	CLEF 2002-03, Q:91-140	225,371	1,938
FR	French	Le Monde 94, SDA French 94-95	CLEF 2002-03, Q:251-350	129,806	3,524

Table 2: Comparison of different dictionary-based short query translation methods. Superscripts indicate that the improvements are statistically significant (2-tail t-test, $p \leq 0.05$). $n-m$ indicates all methods in range $[n, .., m]$.

		FR (short)			DE (short)			ES (short)			IT (short)		
	ID	MAP	P@5	P@10	MAP	P@5	P@10	MAP	P@5	P@10	MAP	P@5	P@10
-	MONO	0.3262	0.4121	0.3737	0.2675	0.4323	0.3688	0.3518	0.4962	0.4321	0.2949	0.3677	0.3115
1	MT	0.2858	0.3939	0.3394	0.2889	0.4375	0.3896	0.3339	0.4280	0.3800	0.2579	0.3510	0.3224
2	TOP-1	0.2211	0.3122	0.2735	0.2015	0.2531	0.2327	0.2749	0.3673	0.3265	0.1566	**0.2208**	0.1896
3	UNIF	0.1944	0.2694	0.2357	0.2148	0.2816	0.2367	0.2362	0.2939	0.2490	0.1526	0.2000	0.1562
4	STRUCT	0.1677	0.25	0.226	0.1492	0.2267	0.2044	0.2472	0.3348	0.3283	0.0994	0.1333	0.1178
5	BiCTM	0.2156	0.3143	0.2755	0.2126	0.2816	0.2612	0.2652^{*}	0.3429	0.3163	0.1504	0.2167	0.1771
6	JCLTRLM	0.1735	0.2687	0.2417	0.1416	0.2178	0.1933	0.2358	0.3522	0.3283	0.1105	0.1733	0.1511
7	MIXWETM	0.2202	0.3143	0.2622	**0.2166**	**0.2166**	**0.2633**	0.2790	0.3755	0.3122	0.1587	0.2125	0.1833
8	CLWETM	$\mathbf{0.2312}^{2-7}$	**0.3306**	**0.2806**	0.2158^{246}	0.2816	0.2551	$\mathbf{0.2915}^{2-7}$	**0.3837**	**0.3367**	$\mathbf{0.163}^{2-7}$	**0.2208**	**0.1937**

4 Experiments

4.1 Experimental Setup

Overview of the used collections is provided in Table 1. The source collection is a pool of Associated Press 1988-89, Los Angeles Times 1994, and Glasgow Herald 1995 collections that are used in previous TREC and CLEF evaluation campaigns for ad-hoc retrieval.

In all experiments, we use the language modeling framework with the KL-divergence retrieval model and Dirichlet smoothing method to estimate the document language models, where we set the dirichlet prior smoothing parameter μ to the typical value of 1000. To improve the retrieval performance, we use the mixture model for pseudo-relevance feedback with the feedback coefficient of 0.5. The number of feedback documents and feedback terms are set to the typical values of 10 and 50, respectively.

All European dictionaries, documents, and queries are normalized and stemmed using the Porter stemmer. Stopword removal is also performed.[1] The Lemur toolkit[2] is employed as the retrieval engine in our experiments.

We use the Google dictionaries in our experiments[3]. In the European languages, we do not transliterate out of vocabulary (OOV) terms of the source languages. The OOVs of the target language are used as their original forms in the source documents, since they are cognate languages. Note that we use the uniform distribution approach as the initial translation model for retrieving top documents. It is worth mentioning that $p(w^t|\mathbf{w}^s)$ in CLTRLM (see Sec. 2) is estimated by a bi-gram coherence translation model (BiCTM) introduced in [9]. Weights of the edges of the graph are estimated by $p(w_j|w_i)$ computed

[1] We use the stopword lists and the normalizing techniques available at `http://members.unine.ch/jacques.savoy/clef/`.

[2] `http://www.lemurproject.org/`

[3] `http://translate.google.com`

Table 3: Comparison of different dictionary-based long query translation methods.

		FR (long)			DE (long)			ES (long)			IT (long)		
	ID	MAP	P@5	P@10	MAP	P@5	P@10	MAP	P@5	P@10	MAP	P@5	P@10
-	MONO	0.4193	0.5354	0.4727	0.3938	0.5280	0.4780	0.5281	0.6720	0.5960	0.3947	0.5022	0.4356
1	MT	0.3395	0.4263	0.3747	0.3436	0.4400	0.4280	0.4208	0.5600	0.478	0.1376	0.1551	0.1306
2	TOP-1	0.3077	0.3960	0.3434	0.2242	0.3080	0.2500	0.3762	0.4800	0.4320	0.2195	0.2800	0.2622
3	UNIF	0.2709	0.3556	0.3091	0.2425	0.2840	0.2540	0.3243	0.3680	0.3340	0.2095	0.2311	0.2000
4	STRUCT	0.1800	0.2646	0.2394	0.2103	0.252	0.2500	0.2951	0.4000	0.3760	0.1942	0.2444	0.2244
5	BiCTM	0.3050	0.3899	0.3505	0.2442	0.3280	0.2780	0.3841	0.4640	0.4340	0.2172	0.2622	0.2422
6	JCLTRLM	0.2266	0.3414	0.2990	0.1520	0.2160	0.1880	0.2734	0.404	0.3500	0.1459	0.2133	0.1756
7	MIXWETM	0.2983	0.3919	0.3485	**0.2652**	0.3400	0.3040	0.3677	0.4280	0.4080	**0.2381**	**0.3022**	**0.2733**
8	CLWETM	**0.3167**$^{2-7}$	**0.4101**	**0.3657**	0.2622$^{2-6}$	**0.3480**	**0.3080**	**0.4029**$^{2-7}$	**0.500**	**0.4620**	0.2380$^{2-6}$	0.2978	0.2667

by SRILM toolkit [4]. BiCTM is also used as p_2 where α is set by 2-fold cross-validation (see Equation 5).

As discussed in Section 3, we used stochastic gradient descent for learning $\mathbf{W}$ which is initialized with random values in $[-1, 1]$; η is set to a small value which also decreases after each iteration. $\mathbf{u}_{w^s}$ and $\mathbf{v}_{w^t}$ are computed based on negative sampling skip-gram introduced in [8]; the size of the window, the number of negative samples, and the size of the vectors are set to typical values of 10, 45, and 50 respectively.

As shown in [1] JCLTRLM outperforms CLTRLM and therefore we opted JCLTRLM as a baseline. The parameters of LDA are set to the typical values $\alpha_d = 0.5$ and $\beta_d = 0.01$. Number of topics in JCLTRLM is obtained by 2-fold cross-validation.

4.2 Performance Comparison and Discussion

In this section we want to compare effectiveness of a number of competitive methods in CLIR. We consider the following dictionary-based CLIR methods to evaluate the proposed method: (*1*) the top-1 translation of each term in the bilingual dictionaries (TOP-1), (*2*) all the possible translations of each term with equal weights (UNIFORM), (*3*) (BiCTM) proposed in [9] , (*4*) the JCLTRLM method proposed in [1], and (*5*) MIXWETM [10]. As bases of comparisons we also provided results of the monolingual runs in each collection (MONO) and Google machine translator (MT). However, our main focus is to investigate superiority of the proposed method compared to the dictionary-based CLIR baselines which are available for almost all pairs of languages.

All the results (on short queries) are summarized in Table 2 and Fig. 1 shows sensitivity of CLWETM to α and the number of feedback documents. As shown in the table, both MIXWETM and CLWETM outperform other methods in terms of MAP, P@5, and P@10 in all the collections. MIXWETM and CLWETM consistently achieved better results compared to others, but CLWETM is clearly more effective than MIXWETM in almost all the datasets. Although CLWETM lost the competition to MIXWETM in DE, but the differences are not statistically significant. One reason for this outcome is the lower performance of BiCTM compared to other collections (see Eq. 5, Table 2, and Fig. 1). Another reason can be the lower sensitivity of the method to n in this collection. As shown in Fig. 1 top-ranked documents in DE are not as helpful as FR, ES, and IT and thus neither CLWETM nor MIXWETM has significant improvements compared to BiCTM.

Although the focus of this research is on dictionary-based CLIR, but it is clear to a reader that in the European collections with short queries the results of MT are higher than all the baselines in dictionary-based CLIR. In the rest of the experiments, we shed light on effectiveness of the methods on verbose queries obtained by concatenating title

[4] http://www.speech.sri.com/projects/srilm/

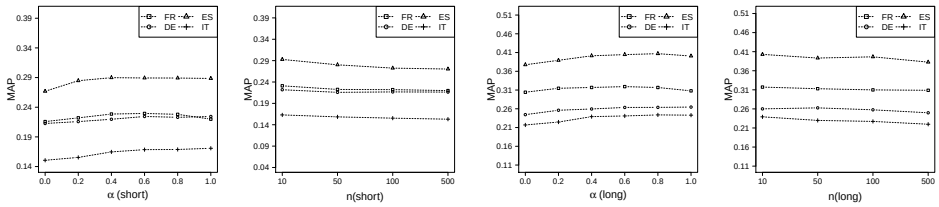

Fig. 1: MAP sensitivity of CLWETM to α and the number of feedback documents.

and description parts of the topics. Table 3 shows the results; the results also confirm the effectiveness of CLWETM over the dictionary-based methods. The most interesting point is decrements of the gaps between CLWETM and MT in quite all the collections. CLWETM reached 93.2%, 76.3%, 95.7%, and 172.9% of the performance of MT in terms of MAP in ES, DE, ES, and IT respectively. In IT we see noticeable decrement of the performance by MT on the verbose queries where the dictionary-based techniques are quite stable.

4.3 Parameter Sensitivity

We investigate the sensitivity of the proposed method to two parameters α and n in Figure 1. We first fix one parameter to its optimal value and then try to get optimal value of the other one . It demonstrates that both parameters work stably across FR, ES and IT collections. The optimal α value empirically is 0.6 and the optimal value of n is 10 in almost all the collections.

5 Conclusion and Future Works

In this paper we presented a translation model for cross-lingual information retrieval, that uses feedback documents in source and target languages for creating word vectors in each one, and then learns a projection matrix to project word vectors in the source language to their translations in the target language. Then we introduced a method for building a translation model that can be easily interpolated with other models using a constant controlling parameter. We investigated the performance of the proposed method on four European collections of CLEF. Our method showed more improvements in FR, SP, and IT since top-ranked documents in DE are not as comparable as the previous ones; therefor applying the proposed method on comparable corpora is considered as an interesting future work. The proposed method reaches up to 87% performance of machine translation (MT) in short queries and considerable improvements in verbose queries.

References

1. Ganguly, D., Leveling, J., Jones, G.: Cross-Lingual Topical Relevance Models. In: COLING '12 (2012)
2. Gouws, S., Bengio, Y., Corrado, G.: Bilbowa: Fast bilingual distributed representations without word alignments. arXiv preprint:1410.2455 (2014)
3. Lavrenko, V., Choquette, M., Croft, W.B.: Cross-lingual relevance models. In: SIGIR '02 (2002)
4. Lavrenko, V., Croft, W.B.: Relevance based language models. In: SIGIR '01. ACM (2001)
5. Lv, Y., Zhai, C.: Positional relevance model for pseudo-relevance feedback. In: SIGIR '10 (2010)
6. Lv, Y., Zhai, C.: Revisiting the divergence minimization feedback model. In: CIKM '14 (2014)
7. Mikolov, T., Le, Q.V., Sutskever, I.: Exploiting similarities among languages for machine translation. arXiv preprint:1309.4168 (2013)
8. Mikolov, T., Sutskever, I., Chen, K., Corrado, G.S., Dean, J.: Distributed representations of words and phrases and their compositionality. In: Adv. in Neur. Inf. Proc. Sys., pp. 3111–3119 (2013)

9. Monz, C., Dorr, B.J.: Iterative Translation Disambiguation for Cross-language Information Retrieval. In: SIGIR '05 (2005)
10. Vulic, I., Moens, M.: Monolingual and cross-lingual information retrieval models based on (bilingual) word embeddings. In: SIGIR '15 (2015)
11. Vulic, I., Smet, W.D., Tang, J., Moens, M.: Probabilistic topic modeling in multilingual settings: An overview of its methodology and applications. IP&M 51(1) (2015)
12. Zamani, H., Dadashkarimi, J., Shakery, A., Croft, W.B.: Pseudo-relevance feedback based on matrix factorization. In: CIKM '16 (2016)